On the scaling limit of finite vertex transitive graphs with large diameter

Itai Benjamini   Hilary Finucane   Romain Tessera

March 1, 2013

Abstract

Let (X_n) be an unbounded sequence of finite, connected, vertex transitive graphs such that $|X_n| = o(\mathrm{diam}(X_n)^q)$ for some $q > 0$. We show that up to taking a subsequence, and after rescaling by the diameter, the sequence (X_n) converges in the Gromov Hausdorff distance to a torus of dimension $< q$, equipped with some invariant Finsler metric. The proof relies on a recent quantitative version of Gromov's theorem on groups with polynomial growth obtained by Breuillard, Green and Tao. If X_n is only roughly transitive and $|X_n| = o\big(\mathrm{diam}(X_n)^\delta\big)$ for $\delta > 1$ sufficiently small, we prove, this time by elementary means, that (X_n) converges to a circle.

1 Introduction

1.1 Scaling limits of transitive graphs

A graph X is *vertex transitive* if for any two vertices u and v in X, there is an automorphism of X mapping u to v. Let (X_n) be a sequence of finite, connected, vertex transitive graphs with bounded degree. Rescale the length of the edges of X_n by the inverse of the graph's diameter so that the resulting metric space X'_n has diameter 1. A metric space $\mathcal{M}$ is the *scaling limit* of (X_n) if (X'_n) converges to $\mathcal{M}$ in the Gromov Hausdorff distance. See e.g. [6, 12] for background on scaling limits and Gromov Hausdorff distance.

In this paper we address the following questions: when is the scaling limit of such a sequence a compact homogeneous manifold? And in that case, what can be said about the limit manifold? By a compact homogeneous manifold, we mean a compact topological manifold M equipped with a geodesic distance (not necessarily Riemannian), such that the isometry group acts transitively on M. Using an old theorem of Turing [22] (see also [13]), one can easily show using standard arguments that any compact homogeneous manifold approximated by finite homogeneous metric spaces must be a torus (see Proposition 2.1.1).

Recall that a Finsler metric on $T^n = \mathbb{R}^n/\mathbb{Z}^n$ is simply a geodesic metric induced by a norm on $\mathbb{R}^n$. When we assume moreover that the degree is bounded, we can be more specific

about the limiting Finsler metric. Let us say that a Finsler metric is polyhedral if the unit ball for the corresponding norm is a polyhedron.

Our main result is the following.

Theorem 1. *Let (X_n) be a sequence of vertex transitive graphs with bounded degree such that $|X_n| \to \infty$ and $|X_n| = o(\mathrm{diam}(X_n)^q)$. Then (X_n) has a subsequence whose scaling limit is a torus of dimension $< q$ equipped with some invariant polyhedral Finsler metric.*

The proof of Theorem 1 makes crucial use of a recent quantitative version of Gromov's theorem obtained by Breuillard, Green and Tao [5], allowing us to reduce the problem from vertex transitive graphs to Cayley graphs of nilpotent groups.

In a former version of this note, we had conjectured that the previous statement should hold without the bounded degree assumption. Since then, Breuillard has provided us with the missing argument (see Section 1.3 for more details).

Theorem 2. *Let (X_n) be a sequence of vertex transitive graphs such that $|X_n| \to \infty$ and $|X_n| = o(\mathrm{diam}(X_n)^q)$. Then (X_n) has a subsequence whose scaling limit is a torus of dimension $< q$ equipped with some invariant Finsler metric.*

Note the slightly weakened conclusion in Theorem 2: indeed, it is not hard to see that any Finsler metric on a Torus of dimension d can be obtained as a scaling limit of a sequence of Cayley graphs (not necessarily with bounded degree) of finite abelian groups.

In the course of proving that the limiting Finsler metric is polyhedral in Theorem 1 we shall prove a result which might be of independent interest: the size of the stabilizers of a faithful transitive action of a step l, rank r nilpotent group on a graph of degree d is in $O_{l,r,d}(1)$ (Lemma 2.3.2).

Let us illustrate Theorem 1 on a simple (although already not completely obvious) example. Given a ring A, one can consider the Heisenberg group $H(A)$ of 3 by 3 upper unipotent matrices with coefficients in A. Now let X_n be the Cayley graphs of the groups $H(\mathbb{Z}/n\mathbb{Z})$ equipped with the finite generating set consisting of the 3 elementary unipotent matrices and their inverses. The cardinality of X_n equals n^3 and easy calculations shows that its diameter is in $\Theta(n)$. For every n, we have a central exact sequence

$$1 \to \mathbb{Z}/n\mathbb{Z} \to H(\mathbb{Z}/n\mathbb{Z}) \to (\mathbb{Z}/n\mathbb{Z})^2 \to 1,$$

whose center is quadratically distorted. In other words, the projection from X_n to the Cayley graph of $(\mathbb{Z}/n\mathbb{Z})^2$ has fibers of diameter $\simeq \sqrt{n}$. It follows that the rescaled sequence (X_n') converges to the 2-torus $\mathbb{R}^2/\mathbb{Z}^2$ equipped with the Finsler metric induced by the ℓ^1-norm on $\mathbb{R}^2$.

1.2 Scaling limits of roughly-transitive graphs

We think an interesting and potentially very challenging open question is to provide a proof of Theorem 2 that does not rely on the results of Breuillard, Green, and Tao. In Section 3, we present an elementary proof of a result similar to the $q = 2$ case of Theorem 2, where

the requirement that X_n be Cayley graphs is weakened, but the assumption on the diameter is strengthened. Recall that for $C \geq 1$ and $K \geq 0$, a (C, K)-quasi-isometry between two metric spaces X and Y is a map $f : X \to Y$ such that

$$C^{-1}d(x, y) - K \leq d(f(x), f(y)) \leq Cd(x, y) + K,$$

and such that every $y \in Y$ is at distance at least K from the range of f. Let us say that a metric space X is (C, K)-roughly transitive if for every pair of points $x, y \in X$ there is a (C, K)-quasi-isometry sending x to y. Let us call a family of metric spaces roughly transitive if there exist some $C \geq 1$ and $K \geq 0$ such that each member of the family is (C, K)-roughly transitive. In Section 3, we will provide an elementary proof of the following theorem.

Theorem 3. *Suppose (X_n) is a roughly transitive sequence of finite graphs such that $|X_n| \to \infty$. There exists a constant $\delta > 1$ such that if*

$$|X_n| = o(\mathrm{diam}(X_n)^\delta),$$

then the scaling limit of (X_n) is S^1.

Modifying slightly the proof of Theorem 3, one can prove that for an infinite, roughly transitive graph X with bounded degree, there exists $R > 0$, $C > 0$ and $\delta > 0$ such that if the volume of a ball of radius $R' \geq R$ is less than $CR'^{1+\delta}$, then X is quasi-isometric to $\mathbb{R}$. Even for vertex-transitive graphs, this provides a new elementary proof (compare [7]).

1.3 Brief outline of the proofs of Theorems 1 and 2

We shall give two (slighty different) proofs of Theorem 1. Our first proof splits into two main steps (ignoring here the explicit bound on the dimension of the limit). First, [5] allows us to reduce the problem to sequences of finite Cayley graphs of virtually nilpotent groups (see Proposition 2.4.1), where the index, step and rank are uniformly bounded. It is important to notice that no bound on the degree of the graphs is required for this step. Then we use the bound on the degree to show that up to taking a subsequence, there exists a fixed infinite virtually nilpotent Cayley graph surjecting to all our finite ones. This allows us to use a result of Pansu [16] to deduce that the limit exists, and is a compact nilpotent (therefore abelian) connected Lie group equipped with some (polyhedral) Carnot-Caratheodory metric.

It is worth mentioning that Pansu's theorem is quite involved and actually not necessary to obtain Theorem 2. On the other hand, Pansu's theorem can be made quantitative [3], so this approach could potentially lead to a quantitative version of Theorem 1.

Another approach is to use the fact that Cayley graphs of nilpotent groups with bounded step and degree enjoy a uniform "doubling property": i.e. for every $r > 0$, a ball of radius $2r$ can be covered by O(1) balls of radius r. This easily implies that the sequence of rescaled graphs are relatively compact for the Gromov-Hausforff topology, and that any accumulation point has finite dimension. Combining this observation with Peter-Weyl's theorem, one obtains that the limit must be a compact nilpotent (hence abelian) connected Lie group.

In the absence of a bound on the degree of our finite graphs, using some version of Pansu's theorem looks definitely trickier than the second approach[1]. The main problem one faces in adapting the second approach is to prove that the X_n are uniformly doubling at all (large enough) scales to ensure convergence. A pigeon-hole principle easily implies that they are doubling at certain (unbounded) scales (see Proposition 2.4.1). What we missed in the previous version of this paper was an argument implying that once such a doubling property occurs at some (large enough) scale, then it "propagates" to larger scales. This turns out to be a consequence of [5] and [4]. We are indepted to Breuillard for explaining it to us.

In both situations, obtaining the bound on the dimension of the limiting torus requires some additional work, including a reduction to the case where the X_n are Cayley graphs of finite abelian groups (see Section 2.6).

Organization: Section 2 is devoted to the proof Theorems 1 and 2. More precisely, Sections 2.1 through 2.3 are devoted to preliminary lemmas. Section 2.4 contains the reduction to virtually nilpotent groups. The proof of Theorem 2 is completed in subsection 2.5, except for the bound on the dimension of the torus. Section 2.6 contains the reduction to abelian groups, allowing us to bound the dimension of the limiting torus in Section 2.7.

The fact that the limiting metric in Theorem 1 is polyhedral naturally follows from the "Pansu's theorem" approach (see Proposition 2.8.1). However, we indicate a way to deduce it from the "Peter-Weyl" approach, only using Pansu's theorem for virtually abelian groups (see Remark 2.8.1).

In Section 3, we prove Theorem 3. In Section 3.2, we provide a second elementary proof of Theorem 3, restoring the assumption that the X_n are vertex transitive, rather than roughly-transitive. Finally in the last section, we discuss the optimality of our results and related open questions.

2 Proof of Theorems 1 and 2

Given a group G with finite generating set S, we will let (G, S) denote the corresponding Cayley graph. Let D_n denote the diameter of X_n. A sequence of groups is said to be uniformly virtually nilpotent if these groups have nilpotent subgroups whose index is bounded above by some constant.

2.1 Preliminaries on Gromov-Hausdorff limits

Proposition 2.1.1. *Let (X_n, d_n) be a sequence of homogeneous finite metric spaces of diameter 1, and for every n, let G_n be a subgroup of isometries acting transitively on X_n. Let*

[1] Such an alternative approach would still be useful in order to obtain a quantitative version of Theorem 2. Here is an idea of how this might be done: by the main result of [5], we know that after dividing by some small normal subgroup and passing to a subgroup of bounded index, balls of intermediate size are "controlled" by homomorphic images of certain "boxes" in some fixed finitely generated free nilpotent group. We are therefore reduced to study the asymptotic shape of large powers of these "box-controlled" sets inside some given finitely generated nilpotent group.

$\hat{d}_n$ be the bi-invariant distance on G_n defined by

$$\hat{d}_n(f,g) = \max_{x \in X_n} d_n(f(x), g(x)).$$

Suppose X_n converges to some compact metric space X. Then $(G_n, \hat{d}_n)$ has a subsequence converging for the GH topology to a closed subgroup $(G, \hat{d})$ of isometries acting transitively on X, where $\hat{d}$ is the bi-invariant distance on G defined by

$$\hat{d}(f,g) = \max_{x \in X} d(f(x), g(x)).$$

Moreover for any sequence $(g_n, h_n) \in G_n^2$ converging to $(g,h) \in G^2$, $g_n h_n^{-1}$ converges to gh^{-1}.

Proof. Since this is standard ([9, 11]), we will only sketch a proof, leaving the details to the reader. The sequence X_n being convergent, it is equi-relatively compact: for all $\epsilon > 0$, there exists N such that X_n is covered by N balls of radius $\epsilon > 0$. It is then easy to check that the sequence $(G_n, \hat{d}_n)$, is also equi-relatively compact, which implies that it is relatively compact for the Gromov-Hausdorff topology.

Now we can argue that the set of accumulation points of a sequence of finite metric Y_n spaces of bounded diameter which is relatively compact for the Gromov-Hausdorff topology coincides with the set of ultralimits of Y_n with respect to all non-principal ultrafilters β on $\mathbb{N}$. More precisely, equip the infinite product $\prod_n Y_n$ with the pseudo-metric $\delta((x_n),(y_n)) = \lim_\beta d(x_n, y_n)$. The ultra-limit of Y_n is then the (compact) metric space obtained by quotienting the previous pseudo-metric space by the obvious equivalence relation.

Now seeing X (resp. G) as the ultralimit of X_n (resp. of G_n) with respect to some ultrafilter β, it is an easy exercise to check that G is naturally a group acting transitively by isometries on X, and that the limiting metric on G is precisely $\hat{d}$, so that G identifies with a subgroup of isometries of X. $\square$

In combination with an old result of Turing [22], one gets the following interesting result.

Corollary 2.1.1. *Let X be a homogeneous riemannian manifold which is the Gromov-Hausdorff limit of a sequence of homogeneous finite metric spaces. Then X is a torus.*

Proof. By the previous proposition, X has a transitive compact (Lie) group of isometries G which is approximable by finite groups in the sense of [22], so that by the main result of [22], the connected component G_0 of G (which also acts transitively) is abelian. It follows that X itself must be a torus. $\square$

2.2 Finite index subgraphs

Lemma 2.2.1. *(Finite index subgraph) Let X be a connected graph of degree d, and G a group acting transitively by isometries on its vertex set. Let $G' < G$ be a subgroup of index $m < \infty$. Let X' be a G'-orbit. Then X' is the vertex set of some G'-invariant graph that is $(O(m), O(m))$-QI to X, and whose degree is bounded by d^{2m+1}.*

Proof. Denote by $[X']_k$ the k-neighborhood of X' in X. Since it is G'-invariant, $[X']_k$ is a union of G'-orbits of X. But X is a union of m G'-orbits, so $X = [X']_m$. Define a G'-invariant graph on X' by adding edges between two vertices of X' if they are at distance at most $2m+1$. Let $y, y' \in X'$ and let $y = x_0, \ldots, x_{v+1} = y'$ be a shortest path between them in the graph X. Let $y_0 = y$ and $y_{v+1} = y'$, and for each $i = 1 \ldots v$, one can find an element y_i in X' at distance at most m from x_i. Clearly the distance between two consecutive y_i is at most $2m+1$, so they are connected by an edge. Thus, $d_{X'}(y, y') \leq d_X(y, y')$. Since $d_X(y, y') \leq (2m+1)d_{X'}(y, y')$, we see that X' is $(2m+1, m)$-QI to X. $\square$

2.3 Preliminary lemmas for (virtually) nilpotent groups

This subsection, culmitating in Lemma 2.3.2, will only be required for the last statement of Proposition 2.4.1, which itself plays a role in subsection 2.4.1. We also believe that Lemma 2.3.2 might be of independent interest. Let us however insist that a reader interested in the proof of Theorem 2 can directly jump to the next subsection.

Let $C^j(G)$ be the descending central series of G, i.e. let $C^0(G) = G$, and $C^{i+1}(G) = [G, C^i(G)]$.

Lemma 2.3.1. *Let G be an l-step nilpotent group, and let S be a symmetric subset of G. Then for every $h \in G$ and every $g \in \langle S \rangle$, $[g, h]$ can be written as a product of iterated commutators of the form $[x_1, [\ldots, x_i] \ldots]$, with $i \leq l-1$, where for each $j = 1 \ldots i$, $x_j \in S \cup \{h^{\pm}\}$, and at least one of the $x_j \in \{h^{\pm}\}$.*

Proof. Without loss of generality we can suppose that G is the free nilpotent group of class l generated by $S \cup \{h^{\pm}\}$. We shall prove the lemma by induction on l. For $l = 1$, the statement is obvious as the group G is abelian. The case $l = 2$ is interesting as it reveals the key idea: one can use the formula $[h, gg'] = [h, g][h, g']$ to break $[h, g]$ into a product of $[h, s]$, where $s \in S$.

Let us assume that the statement is true for $l \leq l_0$ and suppose that G is (l_0+1)-step nilpotent. Applying the induction hypothesis to $G/C^{l_0}(G)$, we deduce that $[g, h]$ is the product of some $y \in C^{l_0}(G)$ with elements $[x_1, [\ldots, x_i] \ldots]$ where $i \leq l_0 - 1$, each $x_j \in S \cup \{h^{\pm}\}$, and at least one x_j in each term is in $\{h^{\pm}\}$.

The element y can be written as a product of elements of the form $[g_1, [\ldots, g_{l_0}] \ldots]$ with $g_j \in G$. Since G is (l_0+1)-step nilpotent, the iterated commutator $[g_1, [\ldots, g_{l_0}] \ldots]$ induces a morphism of abelian groups $\bigotimes_{i=1}^{l_0} G \to C^{l_0}(G)$. Writing each g_j as a word in $S \cup \{h^{\pm}\}$ and using this morphism, we can write y as a product of terms of the form $[a_1, [\ldots, a_{l_0}] \ldots]$ with $a_i \in S \cup \{h^{\pm}\}$.

We have shown that $[g, h]$ can be written as a product of iterated commutators in $S \cup \{h^{\pm}\}$; it remains to show that each term from y contains at least one $x_j \in \{h^{\pm}\}$. The terms obtained from y commute with each other, so we can gather the terms without $h^{\pm}$ into a single word w. Let N denote the normal subgroup generated by h. Since $[g, h] = h^g h^{-1}$, we know that $[g, h] \in N$. Similarly, each iterated commutator containing $h^{\pm}$ is in N, so $w \in N$. But w is also in the subgroup H generated by S, and because G is the free nilpotent group of

class (l_0+1) generated by $S \cup \{h^{\pm}\}$, we have that $N \cap H$ is trivial. Thus w is trivial. This completes the proof of the lemma. $\square$

Corollary 2.3.1. *Let G a l-step nilpotent group generated by some symmetric set S. Then for every element $h \in G$, the normal subgroup generated by h is generated as a subgroup by the elements h^x, where $x \in (S \cup \{h^{\pm}\})^k$, with $k \leq 4^l$.*

Proof. Note that $h^g = [g,h]h$. Applying the lemma to the commutator $[g,h]$ yields a product of iterated commutators with letters in $S \cup \{h^{\pm}\}$, where $h^{\pm}$ appears at least once in each. We leave to the reader to check that such a commutator is a product of conjugates of $h^{\pm}$ by elements x whose word lengths with respect to $S \cup \{h^{\pm}\}$ are at most that of an iterated commutator $[a_1, [\ldots, a_{l-2}]\ldots]$. The length k_l of such commutator is defined inductively as $k_1 = 0$, and $k_{l+1} = 2k_l + 2$. Thus, $k_l \leq 4^l$. $\square$

The following lemma is of independent interest (compare [18]).

Lemma 2.3.2. *Let X be a transitive graph with degree $d < \infty$, and let G be a nilpotent group acting faithfuly and transitively on X. Then the cardinality of any vertex stabilizer H is in $O_{d,r,l}(1)$, where r and l are the rank and step of G. More precisely, let $H = \oplus_p H_p$ be its p-torsion decomposition (which holds since H is nilpotent). Then $\max_p p^n \leq d^{8^l}$, where p^n is the maximal order of an element of H_p.*

Proof. Recall that the rank of any subgroup of G is in $O_{r,l}(1)$, hence the second part of the lemma implies the first one.

The set of vertices of X can be identified with G/H, with x corresponding to the trivial coset H. Faithfulness of the action is equivalent to the fact that H does not contain any non-trivial normal subgroup of G. Observe that the action by left-translation by H preserves the set of neighbors of H. As in the proof of Lemma 2.4.2, we equip G with a generating set S such that X is isomorphic to the Schreier graph $(G/H, S)$, which implies that S is bi-H-invariant. In particular one has $H \subset S^2$.

The proof roughly goes as follows: for every element of H_p of order p^n, we will show that there is a vertex in the ball of radius 8^l of X whose orbit under the action of H has cardinality at least p^n.

By assumption, the normal subgroup generated by $h^{p^{n-1}}$ is not contained in H. Thus by Corollary 2.3.1, there exists $g \in (S \cup H)^{4^l} \subset S^{8^l}$ such that $g^{-1}h^{p^{n-1}}g$ does not belong to H. But then this implies that for all $i = 1, \ldots, p^n$, the vertices $h^i gH$ are distinct, for if $h^i gH = h^j gH$ for some $1 \leq i < j \leq p^n$, then $y = g^{-1}h^{i-j}g \in H$. Write $i - j = p^a b$, with $a < n$ and b coprime to p. There is some c so that $cb = 1 \mod p^n$. Then H contains $y^{cp^{n-1-a}} = g^{-1}h^{p^{n-1}}g$, which is not in H, giving a contradiction. Hence, the $h^i gH$ are distinct for every $1 < i < j \leq p^n$. Since g has length at most 4^l, and the left-translation by H preserves the distance to the origin, the number of such translates is at most the cardinality of the ball of radius 8^l, i.e. at most d^{8^l}, so we are done.

Observe that the same argument can be used to prove that H does not contain any element of infinite order. $\square$

Corollary 2.3.2. *Lemma 2.3.2 extends to virtually nilpotent groups. Namely, let X be a transitive graph with degree $d < \infty$, and let G be a group acting faithfuly and transitively on X. Assume G has a (r,l)-nilpotent subgroup of finite index i. Then the cardinality of any vertex stabilizer H is in $O_{d,r,l,i}(1)$.*

Proof. This easily follows from Lemma 2.2.1 and Lemma 2.3.2. □

2.4 Reduction to virtually nilpotent groups

The first (and main) step of the proofs of both Theorem 1 and Theorem 2 is the following reduction.

Proposition 2.4.1. *With the assumptions of Theorem 2, there exists a sequence of uniformly virtually nilpotent groups N_n with uniformly bounded step and generating sets T_n such that X_n is $(1, o(D_n))$-quasi-isometric to (N_n, T_n). Moreover, if we assume in addition that the degree of X_n is bounded, then we can take T_n with bounded cardinality, and $|N_n|$ to be less than a constant times $|X_n|$.*

Remark 2.4.1. *In the proof of the proposition, N_n appears as a group of isometries, though not of X_n, but rather of a quotient graph Y_n. Precisely we have the following diagram:*

$$\begin{array}{ccc} & & X_n \\ & & \downarrow \\ (N_n, T_n) & \longrightarrow & Y_n \end{array} \tag{1}$$

where the arrows are graph projections, whose fibers have diameter in $o(D_n)$.

Fix some $x \in X_n$. Let G_n be the automorphism group of X_n, H_n the stabilizer of x, and S_n be the subset of G_n containing all g such that $g(x)$ is a neighbor of x. When it is clear from context, let G_n denote the Cayley graph of G_n with generating set S_n. We will need the following lemmas.

Lemma 2.4.1 (Doubling in X_n)**.** *There exists a K depending only on q so that the following holds. For all $R_0 > 0$ there exists an $R = R(n) > R_0$ so that $|B_{X_n}(100R)| \leq K|B_{X_n}(R)|$ for all large enough n, and $R(n) = o(D_n)$.*

Proof. Suppose $|B_{X_n}(100R)| > K|B_{X_n}(R)|$ for all $R_0 < R < D_n^{1/2}$. Then

$$|X_n| \geq |B_{X_n}(D_n^{1/2})| > K^{\log_{100}(D_n^{1/2}/R_0)}|B_{X_n}(R_0)| \geq CD_n^{(1/2)\log_{100} K},$$

for some C independent of n. Letting $K = 100^{q+1}$, there is some N such that for $n > N$ this cannot hold. □

Lemma 2.4.2 (Doubling in G_n)**.** *There exists a K depending only on q so that the following holds. For all $R_0 > 0$ there exists an $R = R(n) > R_0$ so that $|B_{G_n}(100R)| \leq K|B_{G_n}(R)|$ for all large enough n, and $R(n) = o(D_n)$.*

Proof. The vertices of X_n correspond naturally to the cosets of H_n, and X_n is isomorphic to the Schreier graph $(G_n/H_n, S_n)$. The projection mapping (G_n, S_n) to $(G_n/H_n, S_n)$ is a graph homomorphism, as $H_nS_nH_n = S_n$, and so for all $h, h' \in H$ we have $y = zs$ for some $s \in S_n$ if and only if $yh = zh's'$ for some $s' \in S$. The homomorphism from (G_n, S_n) to $(G_n/H_n, S_n)$ sends $B_{G_n}(r)$ to $B_{X_n}(r)$ for all $r \geq 0$. So both sides in Lemma 2.4.1 are multiplied by the same constant. ☐

The main tool in our proof is the following theorem from [5] (although not exactly stated this way in [5], it can be easily deduced from [5, Theorem 1.6] using the arguments of the proof of [5, Theorem 1.3]).

Theorem 2.4.1 (BGT)**.** *Let $K \geq 1$. There is some $n_0 \in \mathbb{N}$, depending on K, such that the following holds. Assume G is a group generated by a finite symmetric set S containing the identity. Let A be a finite subset of G such that $|A^5| \leq K|A|$ and $S^{n_0} \subset A$. Then there is a finite normal subgroup $F \triangleleft G$ and a subgroup $G' \subset G$ containing F such that*

- *G' has index $O_K(1)$ in G*
- *$N = G'/F$ has step and rank $O_K(1)$.*
- *F is contained in $A^{O_K(1)}$.*

Proof of Proposition 2.4.1. By Lemma 2.4.2, we can find a sequence R_n with both R_n and D_n/R_n tending to infinity such that $|B_{G_n}(100R_n)| \leq K|B_{G_n}(R_n)|$. Then by Theorem 2.4.1 applied to $A = B_{G_n}(R_n)$, we obtain a sequence of groups F_n and $M_n = G_n/F_n$ such that F_n has diameter $o(D_n)$, and M_n is uniformly virtually nilpotent with uniformly bounded step and rank.

Proposition 2.4.1 now results from the following facts.

- Let Y_n be the graph obtained by quotienting X_n by the normal subgroup F_n. The quotient map $X_n \to Y_n$ has fibers of diameter $o(D_n)$, so Y_n is $(1, o(D_n)$-QI equivalent to X_n. Note that M_n acts transitively on Y_n, and that if X_n had bounded degree, then Y_n also has bounded degree.
- Let L_n be the kernel of the action of M_n on Y_n, and define $N_n = M_n/L_n$. Then N_n acts faithfuly and transitively on Y_n, so Y_n is isomorphic to $(N_n/H_n, T_n)$, where H_n is the stabilizer of a vertex x, and T_n is the set of elements of N_n taking x to a neighbor of x. As in the proof of Lemma 2.4.2, Y_n is $(1, O(1))$ quasi-isometric to (N_n, T_n).
- (For the second statement of Proposition 2.4.1) Assume here that the degree of X_n is bounded. Recall that N_n is uniformly virtually nilpotent of bounded step and rank and Y_n has bounded degree. Thus, we deduce from Lemma 2.3.2 and Corollary 2.3.2 that H_n has uniformly bounded cardinality, and so does T_n.

Recapitulating, we constructed quasi-isometries between X_n and Y_n, and finally between Y_n and the Cayley graph of N_n. The multiplicative constants of these quasi-isometries are all equal to 1, and the additive ones are bounded by $o(D_n)$, so the proposition is proved. ☐

2.5 The scaling limit in Theorems 1 and 2 is a torus with a Finsler metric

We will use the following result, based on [4] and [5], which was communicated to us by Breuillard (he informed us that he will include a proof of this fact in a subsequent note with his co-authors Green and Tao).

Theorem 2.5.1. *Let G be a group generated by some symmetric finite subset S. For every $K > 0$, there exists $n_0 \in \mathbb{N}$ and $K' > 0$ such that if for some $n \geq n_0$, $|S^{100n}| \leq K|S^n|$, then for all $m \geq n$, one has $|S^{100m}| \leq K'|S^m|$.*

We will now show that the scaling limit in Theorems 1 and 2 is a torus. Recall that in the proof of Lemma 2.4.2, we showed that the isometry group G_n, equipped with some suitable word metric, is such that the volume of its balls of a given radius equals a constant times the volume of balls of same radius in X_n. Therefore, doubling at a single scale in G_n (established in Lemma 2.4.2) implies doubling at any scale in G_n by Theorem 2.5.1, which then implies doubling at any scale in X_n. Thus, the X_n are relatively compact for the Gromov-Hausdorff topology (see for instance [12, Proposition 5.2]). Let X be an accumulation point, and consider now only the subsequence which converges to X. As mentioned in the introduction, the uniform doubling property at any scale also implies that X has finite Hausdorff dimension.

By Proposition 2.4.1, we know that there are virtually nilpotent groups N_n with generating sets T_n such that the Cayley graphs (N_n, T_n) are $(1, o(D_n))-$quasi-isometric to X_n. Thus, X is also a scaling limit of (N_n, T_n). Proposition 2.1.1 allows us to conclude that $(N_n, \hat{d}_n)$ has a subsequential limit N which acts faithfully and transitively on X. Applying the second half of Proposition 2.1.1, because N_n has bounded index, step and rank, N must also be virtually nilpotent.

Observe that the Gromov-Hausdorff distance from X_n to its 1-skeleton is 1, and that the latter is a geodesic metric space. Hence because the scaling limit of X_n is X, its rescaled 1-skeleton also converges to X, which is therefore a geodesic compact metric space. We can now apply a consequence of Peter-Weyl's theorem due to Montgomery and Zippin [15, 1.6.4 p 123]: if a compact group acts transitively and faithfully on a finite dimensional path connected and locally connected topological space, then it is a Lie group. Applying this to N yields that its connected component N^0 must be a connected nilpotent compact Lie group, hence a torus. Now it is easy to deduce that X itself must be a torus on which N^0 acts transitively.

It is not difficult then to see that any invariant geodesic metric on a torus is Finsler. This is equivalent to the fact that any translation-invariant proper geodesic metric d on $\mathbb{R}^m$ is associated to a norm. This statement finally reduces to proving that balls with respect to d are convex (the fact that they are symmetric resulting from translation-invariance). Here is a quick proof of this fact. Fix a euclidean metric d_e on $\mathbb{R}^m$.

Lemma 2.5.1. *Let K be a compact symmetric subset of $\mathbb{R}^m$ containing 0, of euclidean diameter $\leq \varepsilon$, and let $\hat{K}$ be its convex hull. Then for every $n \in \mathbb{N}$, the Gromov-Hausdorff distance relative to the euclidean distance from K^n to its convex hull is at most $(m+1)\varepsilon$.*

Proof. Clearly the convex hull of K^n coincides with $\hat{K}^n$. Now let $x \in \hat{K}^n$, then $x = ny$, where $y \in \hat{K}$. Now y can be written as a convex combination $y = t_0 y_0 + \ldots + t_m y_m$ of $m+1$ elements of K. Write $nt_i = n_i + s_i$, where $s_i \in [0,1)$, and $n_i = [nt_i]$, so that $\sum_i n_i \leq n$ and $\sum_i s_i \leq m+1$. Then, $x = ny = u + z$, where $u \in K^n$, and where $z \in \hat{K}^{m+1}$. This proves the lemma. □

Fix some radius $r > 0$ and some small number $\varepsilon > 0$. Let n be some large enough integer so that $(m+1)r/n \leq \varepsilon$. By the lemma, $B_d(0,r) = B_d(0,r/n)^n$ is at GH-distance at most ε from its convex hull. Since ε was chosen arbitrarily small, this implies that $B(0,r)$ is convex, so we are done.

2.6 Reduction to Cayley graphs of abelian groups

In Section 2.5, we showed that X_n is relatively compact for the Gromov-Hausdorff distance and that accumulation points are finite dimensional tora. Moroever, we proved along the way that there exists a sequence of finite Cayley graphs (N_n, T_n) with the same accumulation points as X_n, such that N_n is virtually nilpotent with bounded index and step. We let $\hat{d}_n$ denote the metric on N_n such that $(N_n, \hat{d}_n)$ acts transitively on (N_n, T_n).

Let N_n^0 be a nilpotent normal subgroup of bounded index in N_n. Let $A_n = N_n/[N_n^0, N_n^0]$, and let A_n^0 be the abelianization of N_n^0. Let d_n be the distance on N_n^0 induced by the word metric on (N_n, T_n), divided by the diameter D_n of (N_n, T_n), and let d_n' be the distance on A_n^0 induced by the word metric on $(A_n, \pi_n(T_n))$ divided by D_n, where $\pi_n : N_n \to A_n$ is the projection map.

Proposition 2.6.1. *Suppose (N_n, d_n) converges to some torus (X, d). Then the sequence of Cayley graphs $(A_n, \pi_n(T_n))$ whose distance is rescaled by D_n, as well as (A_n^0, d_n') have subsequences converging to (X, d).*

Proof. Being at bounded Gromov-Hausdorff distance, these two sequences have the same limit. Let us focus then on (A_n^0, d_n'). Since it has bounded index in N_n, (N_n^0, d_n) converges to (X, d). The projection from (N_n^0, d_n) to (A_n^0, d_n') is 1-lipschitz, so it suffices to show that the fibers have diameter in $o(1)$. By transitivity, this is equivalent to checking that any sequence $g_n \in [N_n^0, N_n^0]$ satisfies $d(1, g_n) \to 0$.

By Proposition 2.1.1, up to passing to a subsequence, $(N_n, \hat{d}_n)$ converges to some compact virtually nilpotent subgroup G of the isometry group of X. Let $\hat{d}_n^0$ denote the metric on N_n^0 induced by $\hat{d}_n$. Again applying the second half of Proposition 2.1.1, we see that $(N_n^0, \hat{d}_n^0)$ converges to a compact nilpotent (hence abelian) subgroup N^0 of N.

By Lemma 2.6.1 (below), g_n can be written as a bounded product of commutators, each of which converges to a commutator in N^0. It follows that $\hat{d}_n^0(1, g_n) \to 0$ for any sequence $g_n \in [N_n^0, N_n^0]$. The proposition then follows from the inequality $d_n \leq \hat{d}_n^0$. □

Lemma 2.6.1. *Let N be some l-step nilpotent group. Then every element x in $[N, N]$ can be written as a product of l commutators.*

Proof. The statement is easy to prove by induction on l. There is nothing to prove if $l = 1$, so let us assume that $l > 1$. Note that $C^l(N) = \{1\}$, so that $C^{l-1}(N)$ is central. Let $x \in [N, N]$. By induction, x can be written as a product of $l-1$ commutators times an element of $C^{l-1}(N)$. Hence it is enough to prove that every element of $C^{l-1}(N)$ can be written as a single commutator. Recall that the iterated commutator $[x_1, [x_2[\ldots, x_l]\ldots]$ induces a morphims from $\bigotimes_{i=1}^{l} A$ to $C^{l-1}(N)$, where $A = N/[N, N]$. In particular its range is a subgroup of $C^{l-1}(N)$. Since it contains generators of $C^{l-1}(N)$ it is equal to $C^{l-1}(N)$, which proves the lemma. $\square$

To complete the reduction, we want to find abelian groups (B_n, V_n) that have diameter of the order of D_n and volume at most that of X_n, that converge to (X, d') where d' is bilipschitz to d, and such that if the degree of X_n is bounded, so is V_n.

To do this, let Y_n be as in Section 2.4. Then we we can extend the diagram from Remark 2.4.1 as follows:

$$\begin{tikzcd} & X_n \arrow[d] \\ (N_n, T_n) \arrow[r] \arrow[d] & Y_n \arrow[d] \\ (A_n, U_n) \arrow[r] & Z_n \\ (A_n^0, U_n^0) \arrow[r] \arrow[u] & Z_n' \arrow[u] \end{tikzcd} \tag{2}$$

where Z_n is the quotient of the Schreier graph Y_n of Remark 2.4.1 by the normal subgroup $[N_n^0, N_n^0]$, Z_n' is the graph guaranteed in Lemma 2.2.1, and the bottom vertical arrows are (O(1),O(1))-QI inclusion as in Lemma 2.2.1.

If the degree of X_n is bounded then the degree of Z_n is bounded. By Lemma 2.2.1, if the degree of Z_n is bounded then the degree of Z_n' is also bounded. Moreover, $|Z_n'| \leq |Z_n| \leq |Y_n| \leq |X_n|$.

The graph Z_n can be described as the Schreier graph of A_n/K_n for some subgroup K_n with respect to the bi-K_n-invariant symmetric subset $K_nU_nK_n$, so that the horizontal arrows from (A_n, U_n) to Z_n and, similarly, from (A_n^0, U_n^0) to Z_n', are graph projections. By the previous proposition the fibers of these projections have diameter in $o(D_n)$, allowing us to conclude both that Z_n has diameter on the order of D_n and that Z_n' has scaling limit (X, d'), where d' is bilipschitz to d.

Since A_n^0 is abelian, the stabilizers of its action on Z_n' are simply its kernel, so that Z_n' itself is the Cayley graph of some abelian group, denoted (B_n, V_n). This concludes our reduction to the case where X_n is the Cayley graph of some abelian group.

2.7 Bound on the dimension of the limiting torus

In this section we prove the bound on the dimension of the limiting torus in Theorems 1 and 2. We shall use repeatedly the following easy fact.

Proposition 2.7.1. *Let (G_n, S_n) be a sequence of Cayley graphs of abelian groups such that $S_n = \{\pm e_1(n), \ldots, \pm e_k(n)\}$ for some fixed $k \in \mathbb{N}$. Then for any sequence Δ_n going to infinity, the rescaled sequence $(G_n, d_{S_n}/\Delta_n)$ subconverges to some connected abelian Lie group of dimension at most k.*

Proof. Let π_n be the projection $\mathbb{Z}^k \to G_n$ mapping the standard generating set S of $\mathbb{Z}^k$ to S_n. Clearly it is enough to show that the Cayley graph $(\mathbb{Z}^k, S)$ has a Δ_n-rescaled limit equal to $\mathbb{R}^k$ (equipped with the standard ℓ^1-metric). This readily follows from Lemma 2.5.1. □

The bound on the dimension in Theorem 1. The bounded degree case (only needed for Theorem 1) being technically easier, we shall prove it first.

Proposition 2.7.2. *Let (B_n, V_n) be a sequence of (finite) Cayley graphs where B_n is abelian, $V_n = \{\pm e_1, \ldots, \pm e_k\}$ for some fixed k. Then up to some subsequence, the rescaled sequence (B_n, V_n) converges to a torus X of dimension at most the largest integer j such that $D_n^j = O(|B_n|)$.*

Before proving the proposition, let us introduce a useful definition.

Definition: Let (G, S) be a Cayley graph where G is abelian, $S = \{\pm e_1, \ldots, \pm e_k\}$. Define its radius of freedom $r_f(G, U)$ to be the largest r such that the natural projection $\mathbb{Z}^k \to G$ is isometric in restriction to the ball of radius r.

Proof of Proposition 2.7.2: Let d denote the dimension of X; we want to show $d \leq j$. Clearly we have $d \leq k$. The strategy is to show that we can reduce the number of generators (changing also the group) to precisely $2d$ such that the limit still has at least the same dimension.

- Claim 1: Let $r_n = r_f(B_n, V_n)$. By definition, $r_n + 1$ is the smallest integer such that there exist $(n_1, \ldots, n_k)$ with $\sum |n_i| \leq 2(r_n + 1)$ such that $n_1 e_1 + \cdots + n_k e_k = 0$. Up to permuting the generators, we can assume that $n_k \neq 0$.
- Claim 2: Let C_n be the subgroup of B_n generated by the set $W_n = \{\pm e_1, \ldots, \pm e_{k-1}\}$. The map $(C_n, W_n) \to (B_n, V_n)$ is a graph morphism and therefore is 1-Lipschitz, and any element in B_n lies at distance at most r_n from its image.
- Claim 3: If $r_n = o(D_n)$, then by Proposition 2.7.1 the sequence of (C_n, W_n) rescaled by D_n subconverges to some connected abelian Lie group which projects continuously onto X.

- Claim 4: Suppose that $r_f(C_n, W_n) = o(D_n)$, then Claim 1 to Claim 3 still hold so that we obtain a sequence (G_n, Y_n) where G_n is the subgroup generated by a symmetric subset Y_n comprising $k-2$ elements (and their inverses) of W_n, and such that a sublimit projects to X.

Iterating this until r_f becomes comparable along some subsequence to D_n, we obtain that for some $l \le k$, and up to reindexing the e_j, the Cayley subgraphs generated by $(\pm e_1, \ldots \pm e_l)$, rescaled by D_n converge to an abelian connected Lie group of dimension at most l: hence $d \le l$. But r_f being of the order of D_n, the volume of this subgraph is at least $\Theta(D_n^l)$. Since it is a subgraph, we have $l \le j$. Hence we duly have $d \le j$.

$\square$

The bound on the dimension in Theorem 2. We shall adapt the previous argument to the case where X_n does not necessarily have bounded degree. Like X_n, (B_n, V_n) enjoys a doubling property at some intermediate scale $R = R(n)$ (i.e. unbounded but small compared to the diameter, see Lemma 2.4.1). According to a celebrated theorem of Freiman and Ruzka [8, 19], this has strong implications on the structure of the ball of radius R.

Definition: Let $\bar{L} = (L_1, \ldots L_k)$ be a sequence of positive integers. The $\bar{L}$-box $B(\bar{L})$ is the subset of $\mathbb{Z}^k$

$$B(\bar{L}) = \{x = \sum_i^k x_i e_i;\ |x_i| \le L_i\}.$$

Given a abelian group $(G, +)$, an $\bar{L}$-arithmetic progression of size $\bar{L}$ of G is a homomorphic image of $B(\bar{L})$.

The following is an easy consequence of Freiman-Ruzsa's theorem as formulated in [20] and [5, Theorem 2.2] .

Theorem 2.7.1. *Let $G = (G, +)$ be an abelian group, and let $K > 2$ be a parameter. Then there exist $q, k \in \mathbb{N}$ depending only on K such that the following holds. Let A be a symmetric finite subset such that $|A^{100}| \le K|A|$. There is a subgroup F of G contained in A^{10}, and a sequence $\bar{L} = (L_1, \ldots L_k)$ with the following properties. Let $G' = G/F$, and A' be the corresponding projection of A. Then $(A')^{10}$ contains an $\bar{L}$-arithmetic progression P and is contained in P^q.*

To allege notation, we shall omit the dependency on n for the following objects: R, L_i (positive numbers), F (subgroup), P and Q (subsets).

As mentioned above, we shall apply this theorem with A being the ball of radius $R = R(n)$ in (B_n, V_n). Note that since the subgroup $F = F(n)$ has diameter in $o(D_n)$, the quotient $(B_n/F, V_n/F)$ still converges to X, so that we can simply assume that F is trivial. Now it is convenient to reformulate the definition of radius of freedom of $P = P(n)$ as follows: $r_f = r_f(n)$ is the largest $r > 0$ such that the homomorphism sending $B(\bar{L})$ to P is injective in restriction to the sub-box of size $(rL_1/R, \ldots, rL_k/R)$.

Observe that R times the word metric with respect to V_n, and the word metric with respect to V_n^R only differ by an additive factor of the order of $R = o(D_n)$. Moreover by Theorem 2.7.1, the word metrics on B_n associated respectively to V_n^R and to P are in fact bi-Lipschitz equivalent (with constants in $O(1)$). It follows that $(B_n, (R/D_n)d_P)$ subconverges to X equipped with some bi-Lipschitz metric.

The proof is then essentially identical to the case of Theorem 1: by the same process we can reduce the number of coordinates of P until the radius of freedom becomes –up to some subsequence– in $\Theta(D_n)$. Strictly speaking, we will not be allowed to apply Proposition 2.7.1 since the generating sets are not finite. However, it is essentially trivial to check that its proof extends to the case when the generating set is a box: more precisely, what we need is that given a sequence of boxes P_n in $\mathbb{Z}^l$ and a positive sequence $\alpha_n \to 0$, the sequence $(\mathbb{Z}^l, \alpha_n d_{P_n})$ subconverges to $\mathbb{R}^l$. As for Proposition 2.7.1 this is an easy consequence of Lemma 2.5.1.

Now let us see what Claims 1 to 3 become in our more general setting. We actually need an additionnal step (say Claim 0) in order to get rid of the potential coordinates such that $L_i = o(R)$. Observe that the projection $\pi_i(V_n^R)$ modulo the subgroup spanned by the other coordinates has cardinality $o(R)$. Let $v \in V_n$ be a generator of the cyclic group $\pi_i(B_n)$. It follows that the order of $\pi_i(v)$ is in $o(R)$. In particular, $\pi_i(B_n)$ has diameter $o(D_n)$, so that $\ker \pi_i$ is at Hausdorff distance $o(D_n)$ from B_n. We can then argue as in Claims 2 and 3 of the proof of Theorem 1 to show that the D_n-rescaled limit of $\ker \pi_i$ surjects to the limit of B_n.

Let us now pass to Claim 1: suppose that $r_f = o(D_n)$, so that there exists a non-zero sequence of integers $(n_1, \ldots, n_k)$ lying in the box $B(3r_fL_1/R, \ldots, 3r_fL_k/R)$ such that $\sum n_ie_i = 0$. Up to permuting the coordinate, assume that n_k is non-zero. Claim 2: we need to check that the subgroup C_n generated by the first $(k-1)$ coordinates is at distance $o(D_n)$ from B_n: by the previous assumption any element in B_n differs from an element in C_n by an element of the form me_k, where $m = o((D_n/R)L_k)$. Such an element has size in $o(D_n/R)$ with respect the to generating set P, so by the remark above, has size $o(D_n)$ with respect to V_n. So we are done with Claim 2. Claim 3 is exactly identical to the case of Theorem 1.

Itterating this procedure, and up to reindexing $\{1, \ldots, k\}$, we end up with some $\bar{L}'$-arithmetic progression Q associated to the box of size $\bar{L}' = (L_1, \ldots, L_l, 0, \ldots, 0)$ for some $l \leq k$, such that the inclusion $(G_n, d_Q) \to (B_n, d_P)$ is 1-Lipschitz, and its range is at Hausdoff distance $o(D_n)$ from B_n, where G_n be the subgroup of B_n generated by Q. Therefore, $(G_n, (R/D_n)d_Q)$ subconverges to some abelian connected Lie group of dimension l, which projects continuously to the scaling limit of X: hence $d \leq l$. Using that none of the remaining L_i's is in $o(R)$, we can therefore conclude that the volume of $Q^{[r_f/R]}$, and therefore of B_n is at least of the order of D_n^l. It follows that the dimension of X is at most the largest integer j such that $D_n^j = O(|B_n|)$.

This ends the proof of the upper bound on the dimension in Theorem 2.

2.8 A more explicit description of convergence in Theorem 1

In this section, we use a theorem of Pansu to give a more explicit description of the convergence to the limiting torus.

Let N be a nipotent connected Lie group, and let $\mathfrak{n}$ be its Lie algebra, and let $\mathfrak{m}$ be a vector supplement of $[\mathfrak{n}, \mathfrak{n}]$ equipped with a norm $\|\cdot\|$. A smooth path $\gamma : [0,1] \to N$ is said to be horizontal if for every t, $\gamma(t) \cdot \gamma'(t)$ belongs to $\mathfrak{m}$ for all $t \in [0,1]$. The length of γ with respect to $\|\cdot\|$ is then defined as

$$l(\gamma) = \int_0^1 \|\gamma(t) \cdot \gamma'(t)\| dt.$$

It can be shown that since $\mathfrak{m}$ generates the Lie algebra $\mathfrak{n}$, every pair of points can be joined by a horizontal path (see [11]). The Carnot-Caratheodory metric associated to $\|\cdot\|$ is defined so that the distance between two points in N is

$$d(x, y) = \inf_{\gamma} \{l(\gamma);\ \gamma(0) = x,\ \gamma(1) = y\},$$

where the infimum is taken over all piecewise horizontal path γ. Note that if $N = \mathbb{R}^k$, so that $\mathfrak{m} = \mathfrak{n}$, then the Carnot-Caratheodory metric is just the distance associated to the norm $\|\cdot\|$, and more generally if N is abelian, d is the Finsler metric obtained by quotienting the normed vector space $(\mathbb{R}^k, \|\cdot\|)$ by some discrete subgroup. Finally we shall need the following simple remark. Suppose $p : N' \to N$ is a surjective homomorphism between nilpotent connected Lie groups, and let d' be a Carnot-Caratheodory metric on N' associated to $(\mathfrak{m}', \|\cdot\|')$. Then there exists a unique metric on N defined so that

$$d(x, y) := \inf_{x', y'} \{d'(x', y');\ p(x') = x,\ p(y') = y\}.$$

This metric is easily seen to be the Carnot-Caratheodory metric on N associated to the normed vector space $(\mathfrak{m}, \|\cdot\|)$ obtained by projecting $(\mathfrak{m}', \|\cdot\|')$.

Proposition 2.8.1. *Suppose (N_n) is a sequence of virtually nilpotent groups with bounded index and step, and bounded generating sets T_n. Then for every $D_n \to \infty$, a subsequence of the Cayley graphs $(N_n, d_{T_n}/D_n)$ converges in Gromov-Hausdorff distance to a connected nilpotent Lie group, equipped with some polyhedral Carnot-Caratheodory metric.*

Note that in Proposition 2.8.1, we do not require D_n to be the diameter of N_n, but rather any unbounded sequence. Specialised in the case where the groups N_n are finite, and D_n is the diameter of the Cayley graph (N_n, T_n), this proposition implies that the limit is a finite dimensional torus (since a compact connected nilpotent Lie group is necessarily a torus), so that the metric on the limit is Finsler.

We will use the following theorem due to Pansu:

Theorem. *[16] Let N be a finitely generated virtually nilpotent group equipped with some finite generating set T, and let $D_n \to \infty$. Then $(N, d_T/D_n)$ converges to some simply connected nilpotent Lie group $N_{\mathbb{R}}$ equipped with some polyhedral Carnot-Caratheodory metric d_{cc}.*

Proof of Proposition 2.8.1. First, up to some subsequence we can assume that the sequence of groups (N_n, T_n) converges to some (N, T). Observe that since N_n is uniformly virtually nilpotent of bounded step, then so is the limit N. Hence N is finitely presented, which implies that for almost all n, (N_n, T_n) is a quotient of (N, T). We shall ignore the finitely many n for which this might fail and denote by $p_n : N \to N_n$ a surjective homomorphism mapping T onto T_n. In particular this implies (N_n, T_n) has a uniform doubling constant. Hence the rescaled sequence is relatively compact for the Gromov-Hausdorff metric [12]. Therefore up to passing to a subsequence, we can suppose that the rescaled sequence converges so some limit space (X, d).

Recall that if a sequence of homogeneous metric spaces (Y_n, d_n) converges to some locally compact space for the Gromov-Hausdorff metric (Y, d), then for any ultra-filter on $\mathbb{N}$, and any sequence of points $o_n \in Y_n$, the corresponding ultra-limit of pointed metric spaces (Y_n, d_n, o_n) is naturally isometric to (Y, d) [12]. Given a sequence $y_n \in Y_n$, let $[y_n]$ denote the equivalence class of (y_n) in the ultralimit.

Let $\tilde{p}_n : (N, d_T/D_n) \to (N_n, d_{T_n}/D_n)$ be the projection induced from p_n. The maps $\tilde{p}_n$ are 1-Lipschitz and surjective, so there is a projection $p : (N_{\mathbb{R}}, d') \to (X, d)$ from the limit $(N_{\mathbb{R}}, d')$ of the sequence $(N, d_T/D_n)$ to (X, d) such that for each sequence (x_n) in N,

$$[\tilde{p}_n(x_n)] = p([x_n]).$$

We also have that for every $g \in N_n$, there is a $x \in N$ so that $p_n(x) = g$ and $|g|_{T_n} = |x|_T$. We claim that X is naturally a group. More precisely if $dist(g_n, g'_n) \to 0$ and $dist(h_n, h'_n) \to 0$ then $[g_n h_n^{-1}] = [g'_n h'^{-1}_n]$. We can write $g'_n = g_n a_n$ and $h'_n = h_n b_n$ with $|a_n|, |b_n| \to 0$. We can choose $x_n, c_n, y_n, d_n \in (N, d_T/D_n)$ that are mapped by p_n to g_n, a_n, h_n, and b_n, respectively, and so that $|c_n|, |d_n| \to 0$ in $(N, d_T/D_n)$. Because $(N, d_T/D_n)$ is a group, we have $[g_n h_n^{-1}] = [p(x_n)p(y_n)^{-1}] = [p(x_n y_n^{-1})] = p[x_n y_n^{-1}] = p([x_n c_n d_n^{-1} y_n^{-1}]) = [p(x_n c_n d_n^{-1} y_n^{-1})] = [g_n a_n b_n^{-1} h_n^{-1}] = [g'_n h'^{-1}_n]$. ∎

Corollary 2.8.1. *The limiting metric in Theorem 1 is polyhedral Finsler.*

Remark 2.8.1. *Let us indicate a shortcut in the proof of Corollary 2.8.1, using only a very simple version of Pansu's theorem. In Section 2.6, we showed that one can assume X is the scaling limit of a sequence of Cayley graphs of virtually abelian groups[2]. Moreover if X_n has bounded degree, then the same holds for these Cayley graphs. (Recall that this followed from Lemma 2.2.1 and did not require the last statement of Proposition 2.4.1, or, therefore, any of Section 2.3.) This means that we can proceed exactly as we did in this subsection, replacing everywhere "virtually nilpotent" by "virtually abelian." In particular Pansu's theorem becomes an exercice in this setting.*

[2]We actually reduced further to abelian groups, but the step of taking a subgroup of finite index involves quasi-isometries with non-trivial multiplicative constants, hence changing the limiting metric (up to bi-Lipschitz homeomorphism).

3 Proof of Theorem 3

In this section, we will provide two elementary proofs of Theorem 3 (the second holding under the stronger assumption that the graphs are transitive).

3.1 First proof

The proof of Theorem 3 goes approximately as follows. We show using rough transitivity that if the graph does not converge to a circle, then it must contain a caret of size proportional to the diameter. Then, iterating using rough transitivity, we generate large volume, contradicting the assumption.

First, we will define a few key terms.

Definition: Given $K \geq 0$ and $C \geq 1$, a (C, K)-*quasi-geodesic* in a metric space X is a (C, K)-quasi isometrically embedded copy of the interval $[1, k]$ into X; i.e. a sequence of points $x_1, \dots x_k \in X$ such that

$$C^{-1}(j-i) - K \leq d(x_i, x_j) \leq C(j-i) + K$$

for all $1 \leq i < j \leq k$.

Let $B(v, r)$ denote the ball of radius r around a vertex v. A quasi-caret of radius $\geq R$ consists of three quasi-geodesic segments γ_1, γ_2 and γ_3 that start from a point v_0, escape from $B(v_0, R)$, and move away from one another at "linear speed." In other words,

Definition: A *quasi-caret of radius R* is a triple γ_1, γ_2, γ_3 of quasi-geodesics from a vertex v_0 to vertices v_1, v_2, and v_3, respectively, such that $d(v_0, v_i) = R$ for $i = 1, 2, 3$, and there is a constant $c > 0$ satisfying that for all k_1, k_2, k_3, $d(\gamma_i(k_i), \gamma_j(k_j)) \geq c \max\{k_i, k_j\}$ for $i \neq j$.

If a roughly transitive graph has a quasi-caret of radius $\geq \epsilon D$, then by moving around this caret with quasi-isometries (with uniform constants), we obtain at every point of the graph a quasi-caret (with uniform constants) of radius $\geq \epsilon' D$.

The proof of Theorem 3 follows from the following four lemmas.

Lemma 3.1.1. *Let $D = \mathrm{diam}(X)$. Suppose there exists a quasi-caret of radius $R = \epsilon D$ for some $\epsilon > 0$ in a finite (C,K)-roughly transitive graph X. Then $|X| \geq \epsilon' D^{\delta}$, where ϵ' and $\delta > 1$ depend only on C, K and ϵ.*

Proof. To avoid complicated expressions that would hide the key idea, we will remain at a rather qualitative level of description, leaving most calculations to the reader.

Given a quasi-caret $(\gamma_1, \gamma_2, \gamma_3)$, we can stack a sequence of disjoint consecutive balls along the γ_j's, whose radii increase linearly with the distance to the center v_0. More precisely, one can find for every $j = 1, 2, 3$ a sequence of balls $B_k^j = B(\gamma_j(i_k), r_k^j)$ such that

- $C^{-1} d(v_0, \gamma(i_k)) \leq r_k^j \leq C d(v_0, \gamma(i_k))$ for some constant $C \geq 1$,
- $r_k^j \geq c' R$ for some $0 < c' < 1$ independent of R,

- $\sum_{j,k} r_k^j \geq \alpha R$ where $\alpha > 1$ is also independent of R,
- the distance between B_k^j and B_{k+1}^j equals 1, and all these balls are disjoint (when j and k vary).

Now, fix some $R \leq \epsilon D$, and consider a ball $B(x, R)$. It contains a quasi-caret of radius R. This caret can be replaced by the balls described above, each one of them containing a quasi-caret of radius r_k^j. The sum of the radii of these carets is at least αR. We can iterate this procedure within each ball, so that at the k-th iteration, we obtain a set of disjoint balls in $B(x, R)$ whose radii sum to at least $\alpha^k R$. Because the radius R decreases by a factor no smaller than c' each time, we can iterate $\log_{1/c'} R$ times. After $\log_{1/c'} R$ iterations, each ball still has positive radius, so $|B(x, R)| \geq \alpha^{\log_{1/c'} R} R = R^{1+\log_{1/c'}(\alpha)}$. Since $\log_{1/c'}(\alpha) > 0$, this proves the lemma. □

Next we will show that under the assumption that no such caret exists, our graphs locally converge to a line. More precisely,

Lemma 3.1.2. *Let X_n be a roughly transitive sequence of graphs of diameter D_n going to infinity whose carets are of length $o(D_n)$. Then there exists $c > 0$ such that for n large enough, any ball of radius cD_n is contained in the $o(D_n)$-neighborhood of a two-sided geodesic line.*

Proof. By rough transitivity, it is enough to prove the lemma for some specific ball of radius cD_n. Start with a geodesic $[x, y]$ of length equal to the diameter D_n. Let z be the middle of this geodesic. We are going to show that the ball of radius $D_n/10$ around z is contained in a $o(D_n)$-neighborhood of $[x, y]$. If this was untrue, we would find a constant c' such that $B(z, D_n/10)$ contains an element w at distance at least $c'D_n$ of $[x, y]$. Now pick an element z' in $[x, y]$ minimizing the distance from w to $[x, y]$. The shortest path from z' to w, together with the two segments of the geodesic starting from z' form a caret of size proportional to D_n. □

If we knew that X_n converges, and that the limit is homogenous and compact, then this lemma would show that the limit is a locally a line, and thus must be S^1. If we knew that there was a large geodesic cycle in X_n, then Lemma 3.1.1 would show that all vertices are close to the cycle, which would also imply that the limit is S^1. However, we know neither of these two facts a priori, so the next lemma is necessary to complete the poof.

Lemma 3.1.3. *Suppose X_n has the property that for some $c > 0$, and for n large enough, any ball of radius cD_n is contained in the $o(D_n)$-neighbourhood of a two-sided geodesic line. Then its scaling limit is S^1.*

Proof. Let $x_1, ..., x_k$ be a maximal $cD_n/10$-separated set of points of X_n, and let B_j be the corresponding balls of radius $cD_n/100$.

We consider the graph H_n whose vertices are labeled by the balls B_j and such that two vertices are connected by an edge if the corresponding balls are connected by a path avoiding the other balls.

By maximality, for any $v \in X_n$, there is at least on x_i in $B(v, cD_n/5)$. Let us consider a fixed x_j. In $B(x_j, cD_n)$, X_n is well approximated by a line, so there are vertices v_1 and v_2 on either side of x_j such that the balls of radius $cD_n/5$ around v_1, v_2, and x_j are disjoint. Thus, there must be an x_i on either side of x_j. Picking on each side the x_i that is closest to x_j, we see that the corresponding balls are connected in H_n. Moreover since removing these two balls disconnects the ball $B(x_j, cD_n/100)$ from all other x_i's, we see that the degree of H_n is exactly two. Hence the graph H_n is a cycle that we will now denote by $\mathbb{Z}/k\mathbb{Z}$. To simplify notation, let us reindex the balls B_j accordingly by $\mathbb{Z}/k\mathbb{Z}$.

Let us show that the graph X_n admits a (simplicial) projection onto $\mathbb{Z}/k\mathbb{Z}$. If we remove the balls B_i, we end up with a disjoint union of graphs $C_1, \ldots, C_k$, such that C_j is connected to B_j and B_{j+1}. Let $V_j = B_j \cup C_j$. The graph V_j connects to and only to V_{j-1} and V_{j+1}. Hence we have a projection from X_n to the cyclic graph $\mathbb{Z}/k\mathbb{Z}$ sending V_j to the vertex j, and edges between V_j and V_{j+1} to the unique edge between j and $j+1$.

Recall that the distance between two consecutive balls B_i and B_{i+1} is at least $cD_n/20$. Now take a shortest loop $\gamma = (\gamma(1), \ldots, \gamma(m) = \gamma(0))$ in X_n among those projecting to homotopically non trivial loops in the graph associated to $\mathbb{Z}/k\mathbb{Z}$. Clearly this loop has length at least $ckD_n/20$ (since it passes through all balls B_i). We claim moreover that it is a geodesic loop. Without loss of generality, we can suppose that $\gamma(0)$ starts in B_0 and that the next ball visited by γ after B_0 is B_1. Observe that although γ might exit some B_i and then come back to it without visiting any other B_j in the meantime, it *cannot* visit B_i, then go to B_{i+1}, and then back to B_i (without visiting other balls in the meantime). Indeed such a bactrack path could be replaced by a shorter path staying within B_i, contradicting minimality. It follows that the sequence of B_i's visited by γ (neglecting possible repetitions) is given by $B_0, B_1 \ldots B_k = B_0$; namely it corresponds to the standard cycle in $\mathbb{Z}/k\mathbb{Z}$. The same argument implies that the sequence of balls visited by a any geodesic joining two points in X_n corresponds to a (possibly empty) interval in $\mathbb{Z}/k\mathbb{Z}$.

Now, suppose for sake of contradiction that γ is not geodesic. This means that there exists an interval of length $\leq m/2$ in γ which does not minimize the distance between its endpoints. But then applying the previous remark, we see that replacing either this interval or its complement by a minimizing geodesic yields a loop whose projection is homotopically non-trivial, hence contradicting our minimal assumption on γ.

We therefore obtain a geodesic loop in X_n whose Hausdorff distance to X_n is in $o(D_n)$. Hence the scaling limit of X_n exists and is isometric to S^1. $\square$

3.2 A second elementary proof

In this section, we present a second elementary proof of Theorem 3, but under the stronger assumption that the X_n are vertex transitive. This proof strategy gives us $\delta = 1 - \frac{1}{\log_3(4)}$.

Theorem 3.2.1. *Suppose X_n are vertex transitive graphs with $|X_n| \to \infty$ and*

$$|X_n| = o(\mathrm{diam}(X_n)^{2-\frac{1}{\log_3(4)}}.$$

Then the scaling limit of (X_n) is S^1.

Let us first describe the rough idea of the proof. To prove Theorem 3.2.1, It suffices to show that in a finite vertex transitive graph with small volume relative to its diameter, there is a geodesic cycle whose length is polynomial in the diameter. If a long caret is rooted at a vertex on this cycle, then using transitivity and iteration we generate large volume, contradicting the assumption. Thus, all vertices must be close to the cycle.

To find a large geodesic cycle, we use the fact that a finite vertex transitive graph X contains a $\mathrm{diam}(X)/8$-fat triangle. If this triangle is homotopic to a point after filling in small faces, this will imply large area and will violate our assumption. Thus, there is a loop that is not contractable. The smallest non-contractable loop is a geodesic cycle, and since we filled in all small cycles, this geodesic cycle must be large.

We begin by proving a version of Lemma 3.1.1 for vertex transitive graphs.

Definition: A *3-caret of branch-length R* is a triple γ_1, γ_2, γ_3 of geodesics from a vertex v_0 to vertices v_1, v_2, and v_3, respectively, such that $d(v_0, v_i) = R$ for $i = 1, 2, 3$, and for all k_1, k_2, k_3, $d(\gamma_i(k_i), \gamma_j(k_j)) \geq \max\{k_i, k_j\}$ for $i \neq j$.

Lemma 3.2.1. *Let $D = \mathrm{diam}(X)$. Suppose there exists a 3-caret of branch length $R = \epsilon D^c$ for some $\epsilon, c > 0$ in a finite vertex transitive graph X. Then $|X| > \epsilon' D^{1+c(\log_3(4)-1)}$, where $\epsilon' = (1/2)\epsilon^{\log_3(4)-1}$.*

Proof. Suppose γ_1, γ_2, and γ_3 form a 3-caret of branch length R. Let u_1, u_2, and u_3 denote the vertices at distance $2R/3$ from v_0 on γ_1, γ_2, and γ_3, respectively, and let $u_0 := v_0$. The u_i are at pairwise distance $2R/3$ from each other, and so $B(u_i, R/3)$ are pairwise disjoint. By vertex transitivity, there is a 3-caret of branch length R centered at each u_i, which intersects $B(u_i, R/3)$ as a 3-caret of branch length $R/3$. Thus, we have four disjoint balls of radius $R/3$, each containing a 3-caret of radius $R/3$.

We can iterate this procedure, dividing R by three at each step and multiplying the number of disjoint balls by four. So for any m, $B(v_0, R)$ contains 4^m balls, each of which contains a 3-caret of branch length $R/3^m$. Letting $m = \log_3(R)$, we have that $B(v_0, R)$ contains 4^m disjoint 3-carets of branch length 1. In particular, $|B(v_0, R)| \geq 4^m = R^{\log_3(4)}$.

There exists a geodesic path γ in X of length D. Let $R = \epsilon D^c$. Then it is possible to take vertices $v_1, \ldots, v_{D/2R}$ in γ such that $B(v_i, R) \cap B(v_j, R) = \emptyset$ for all $i \neq j$. Summing the number of vertices in $B(v_i, R)$ for $1 \leq i \leq D/(2R)$, and using that $|B(v_i, R)| \geq R^{\log_3(4)}$, we have

$$|X| \geq D/(2R) \cdot R^{\log_3(4)} = (1/2)\epsilon^{\log_3(4)-1} D^{1+c(\log_3(4)-1)}.$$

□

The fact that a 3-caret of branch length R implies $|B(v_0, R)| \geq R^{\log_3(4)}$ also has consequences for infinite vertex transitive graphs. For example, if an infinite vertex transitive graph X has linear growth, then there is an upper bound on the size of a 3-caret in X. Since X has a bi-infinite geodesic γ and a vertex at distance R from γ implies a 3-caret of branch length R, every vertex in X must be within a bounded neighborhood of γ. Conversely, if

X does not have linear growth, then there must be vertices at arbitrary distances from any fixed bi-infinite geodesic. Thus X must have growth at least $O(n^{\log_3(4)})$.

Next, we will show that every vertex-transitive graph with a large diameter has a large geodesic cycle. We will use the following theorem from [1].

Definition: A geodesic triangle with sides s_1, s_2, s_3 is δ-fat if for every vertex v in X,

$$dist(v, s_1) + dist(v, s_2) + dist(v, s_3) \geq \delta.$$

Theorem 3.2.2. *Every finite vertex transitive graph with diameter D contains a $(1/8)D$-fat triangle.*

For completeness, here is the short proof. Given vertices u and v, let uv denote a shortest path from u to v

Proof. Suppose X is finite and transitive, and D is its diameter. Let w and z realize the diameter, i.e. $|wz| = D$. By transitivity there is a geodesic path xy that has z as its midpoint and length D. Suppose the triangle wxy is not δ-fat. Then there is a point a on xy such that the distance from a to wy is at most 2δ and the distance from a to wx is at most 2δ. Suppose, w.l.o.g. that a is closer to x than to y. We have $|ax| + |ay| = D$, $|wa| + |ax| < 2\delta + D$ (because a is within 2δ of wx), $|wa| + |ay| < 2\delta + D$. Add these latter two and subtract the previous equality, and get $|wa| < D/2 + 2\delta$. Since $|wz| = D$, this means that $|za| > D/2 - 2\delta$. Since a is on xy and closer to x, this means that $|xa| < 2\delta$. Since a is within 2δ from wy, we have $|wy| > |wa| + |ay| - 2\delta$. Since $|xa| < 2\delta$ and $|xy| = D$ this gives $|wy| > |wa| + D - 4\delta$. Since $|wy|$ is at most D, this implies $|wa| < 4\delta$. But $|za|$ is at most $D/2$. so $D = |wz| \leq |wa| + |za| < 4\delta + D/2$ So $D < 8\delta$. ☐

Lemma 3.2.2. *Suppose X is a finite d-regular vertex-transitive graph such that $|X| < (\alpha/d)D^{2-c}$, where $\alpha = \sqrt{3}/576$. Then X contains a geodesic cycle of length D^c.*

Proof. We will begin by proving two claims.

Claim 1: Suppose H is a d-regular planar graph, every face of H except the outer face has a boundary of length at most D^c, and H contains a $(1/8)D$-fat geodesic triangle. Then $|H| > (\alpha/d)D^{2-c}$.

This is a variant of Besicovich' lemma for squares. Fill each face f of H with a simply connected surface of area at most $|f|^2$ so that distances in H are preserved (for example, a large portion of a sphere), and consider the $(1/8)D$-fat geodesic triangle in the simply connected surface X obtained. The triangle has sides s_1, s_2, and s_3, of lengths at least $(1/8)D$. The map f from X to $\mathbb{R}^3_+$ taking a point x to $(dist(x, s_1), dist(x, s_2), dist(x, s_3))$ is 3-Lipschitz, so the area of the image of f is smaller than 9 times the area of X. For each (x_1, x_2, x_3) in the image we have $x_1 + x_2 + x_3 > (1/8)D$, so projecting radially to the simplex $x_1 + x_2 + x_3 = (1/8)D$ does not increase the area. The projection of the image of the boundary of the triangle is the boundary of the simplex, so the projection is onto. Thus, the area of X is bigger than $1/9\sqrt{3}((1/8)D)^2$.

Each face f of H contributes $|f|^2$ to the area to X, so we can say that each vertex on the border of f contributes $|f|$ to the area of X. Each vertex of H participates in at most d faces, each of which is of size at most D^c so area$(X) \leq dD^c|H|$. Thus $|H| > \alpha/dD^{2-c}$ where $\alpha = (1/9)\sqrt{3}(1/8)^2$.

Claim 2: Suppose X is a finite d-regular graph that contains a $(1/8)D$-fat triangle A, and let T denote the topological space obtained from X by replacing each cycle of length at most D^c with a euclidean disc whose boundary matches the cycle. If A is homotopic in T to a point, then $|X| > (\alpha/d)D^{2-c}$.

Suppose A is homotopic to a point. Then a continuous map from S^1 to A can be extended to a continuous map from the disk B^1 to $T(X, D^c)$. The image of this map has a planar sub-surface S with boundary A. Intersecting S with X, we obtain a planar subgraph H of X such that each face has a boundary of length at most D^c except for the outer face, which has A as a boundary. By Claim 1, $|X| \geq |H| > (\alpha/d)D^{2-c}$.

Now we will prove the lemma. Because $|X| < (\alpha/d)D^{2-c}$, Claim 2 tells us that T is not simply connected. Any loop in T is homotopic to a loop in X, so since T is not simply connected, there exists a topologically non-trivial loop in X. Let ℓ denote the non-trivial loop in X which has minimal length.

Given any two vertices u and v in ℓ, the shortest path from u to v is homotopic to at most one of the two paths p_0, p_1 in ℓ between u and v; say it is not homotopic to p_0. If the length of p^* is shorter than the minimum length of p_0 and p_1, then it would be possible to replace p_1 with p^* to obtain a loop ℓ' which is non-trivial and shorter than ℓ. Thus, the shortest path in X between any two vertices in ℓ is a path in ℓ, so ℓ is a geodesic cycle.

Because all cycles of length less than D^c are homotopic to a point in T, ℓ must have length greater than D^c. □

Proof of Theorem 3.2.1. Let D_n denote the diameter of X_n and $c = \frac{1}{\log_3(4)}$. It suffices to show that for large enough n, there is a geodesic cycle C_n in X_n such that $|C_n| > {D_n}^c$ and

$$\max_{v \in X_n}(dist(v, C_n)) = o({D_n}^c).$$

For large enough n, $|X_n| < (\alpha/d){D_n}^{2-c}$, so Lemma 3.2.2 guarantees the existence of a geodesic cycle C_n with $|C_n| > {D_n}^c$. By Lemma 3.2.1, if there were a 3-caret of branch length $\epsilon {D_n}^c$, we would have $|X_n| \geq \epsilon' {D_n}^{1+c(\log_3(4)-1)} = \epsilon' {D_n}^{2-c}$. So for all ϵ and large enough n, there is no 3-caret of branch length $\epsilon {D_n}^c$. But for $\epsilon < {D_n}^c/4$, a vertex at distance $\epsilon {D_n}^c$ from C_n implies a 3-caret of branch length $\epsilon {D_n}^c$. Thus, for every ϵ and for large enough n, all vertices in X_n are within distance $\epsilon {D_n}^c$ from C_n. □

3.3 What about an analogue of Theorem 3 in higher dimensions?

Theorem 3 cannot be generalized to "higher" dimensions (as in Theorem 2) because any compact manifold can be approximated by a roughly transitive sequence of graphs. Moreover,

there are sequence of roughly transitive graphs with no convergent subsequence, but with a good control on the volume. Here we will sketch the construction of such a sequence.

Recall that compactness is closed under Gromov-Hausdorff limit. Hence a limit X of our sequence X_n, if a limit exists, is necessarily compact. As a result, X has doubling property at any fixed scale: in particular there exists $k \in \mathbb{N}$ such that any ball of radius $\operatorname{diam}(X)/2$ is covered by k balls of radius $\operatorname{diam}(X)/4$. Then for n large enough, balls of radius $\operatorname{diam}(X_n)/2$ are covered by $2k$ balls of radius $\operatorname{diam}(X)/4$.

We start by picking a sequence of Cayley graphs Y_n with no converging subsequence (but without control on the volume). For example, let S be a finite generating subset of $G = \mathrm{SL}(3, \mathbb{Z})$, and let Y_n be the Cayley graph of $G_n = \mathrm{SL}(3, \mathbb{Z}/n\mathbb{Z})$ associated to the (projected) generated set S. The fact that Y_n is an expander violates the previous doubling condition, and hence Y_n does not have any converging subsequence (we leave this easy and standard fact to the reader).

Then convert G_n-equivariantly the Y_n into Riemannian surfaces S_n by replacing edges with empty tubes, and smoothing the joints that correspond to vertices in Y_n. The radius of the tubes in S_n will be some L_n to be determined later, and the length will be $2L_n$. Independently of the choice of L_n, these S_n are uniformly roughly transitive. This follows because the S_n are all rescaled versions of manifolds S'_n that cover the same compact manifold M. For any two points $x, y \in M$, there is a diffeomorphism of M taking x to y, with uniform bounds on the derivative and the derivative of the inverse. This property extends to the covering manifolds S'_n, and thus to S_n.

To obtain the sequence X_n of roughly transitive graphs, replace S_n with a tiling with bounded faces, chosen so that the X_n remain roughly transitive. Choose L_n large enough so that $|X_n| = o(\log(\operatorname{diam}(X_n)) \operatorname{diam}(X_n)^2)$. This is possible, for example, by choosing L_n to grow asymptotically faster than $\frac{1}{D_n^2}(|Y_n| - D_n^2 \log D_n)$, where D_n denotes the diameter of Y_n. Let X'_n denote X_n normalized by $\operatorname{diam}(X_n)$. It remains to show that X'_n has no convergent subsequence. However, the Gromov-Hausdorff distance from X'_n to Y_n is roughly $L_n/\operatorname{diam}(X_n) = 1/\operatorname{diam}(Y_n) \to 0$. Since Y_n has no convergent subsequence, neither does X'_n.

Remark 3.3.1. *Note that in the above construction we can bound the volume by a function of the diameter which is as close as we want to quadratic. But this leaves open the problem of finding a sequence of roughly transitive graphs with subquadratic growth which does not admit a converging subsequence. Also we do not know what could be the best δ for which Theorem 3 holds ($\delta = 2$?). Finally finding a converging counter-example to Theorem 3 with $\delta < 2$ would be of special interest as the limit would be quite an exotic object: it would be a compact geodesic metric space with Hausdorff dimension in $(1, 2)$, and such that for every pair of points $x, y \in X$, there exists a C-bilipschitz homeomorphism sending x to y, where C only depends on X. We do not know if such object exists.*

Acknowledgement: We are grateful to Pierre Pansu for useful discussions and especially to Emmanuel Breuillard for directing us to Theorem 2.4.1 and also for his later interventions (see the introduction). We would also like to thank Tsachik Gelander for pointing us to

Turing's theorem, and its applications to homogeneous manifolds. Last but not least we would like to thank László Pyber for noticing a mistake in a previous version of Lemma 2.3.2. He also mentioned to us some work in progress where some form of Lemma 2.3.2 is proved.

References

[1] I. Benjamini and O. Schramm, Finite transitive graph embeddings into a hyperbolic metric space must stretch or squeeze. *GAFA seminar*, Springer LNM (2012), To appear.

[2] E. Breuillard, Geometry of groups of polynomial growth and shape of large balls. Preprint.

[3] E. Breuillard and Enrico Le Donne. On the rate of convergence for nilpotent groups and subFinsler geometry. Preprint.

[4] E. Breuillard, B. Green, Approximate groups I : the torsion free nilpotent case. J. de lInstitut de Math. de Jussieu 2011.

[5] E. Breuillard, B. Green and T. Tao, The structure of approximate groups. arXiv:1110.5008

[6] D. Burago, Y. Burago and S. Ivanov, *A Course in Metric Geometry*. American Mathematical Society (2001).

[7] M. DeVos and B. Mohar. Small separations in vertex-transitive graphs. *Electronic Notes in Discrete Mathematics* 24: 165-172 (2006).

[8] G. A. Freiman. Foundations of a structural theory of set addition, Translated from the Russian. Translations of Mathematical Monographs, Vol 37. American Mathematical Society, Providence, R. I. (1973), vii+108 pp.

[9] K. Fukaya, T. Yamaguchi. The fundamental groups of almost nonnegatively curved manifolds. Annals of Mathematics, 136 (1992), 253–333.

[10] M. Gromov. Groups of polynomial growth and expanding maps. *Publ. Math. IHES*, 53, 53-73 (1981).

[11] M. Gromov, Carnot-Carathéodory spaces seen from within, in Sub-Riemannian Geometry, edited by A. Bellaiche and J-J. Risler, 79-323, Birkauser (1996).

[12] M. Gromov. Metric structures for Riemannian and non-Riemannian spaces, volume 152 of Progress in Mathematics. Birkhäuser Boston Inc., Boston, MA, 1999. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes.

[13] D. Kazhdan, On ε-representation. Israel J. Math., 43 4 (1982). 315–323.

[14] B. Kleiner. A new proof of Gromov's theorem on groups of polynomial growth. *Jour. AMS* 23, 815-829 (2010).

[15] D. Montgomery and L. Zippin. Topological transformation groups. Interscience Publishers, New York-London, 1955.

[16] P. Pansu. Croissance des boules et des géodésiques fermées dans les nilvariétés. [J] Ergodic Theory Dyn. Syst. 3, 415-445 (1983).

[17] F. Peter, H. Weyl, Die Vollstndigkeit der primitiven Darstellungen einer geschlossenen kontinuierlichen Gruppe, Math. Ann. 97 (1927), 737–755.

[18] C. E. Praeger, L. Pyber, P. Spiga and E. Szabó. The Weiss conjecture for locally primitive graphs with automorphism groups admitting composition factors of bounded rank. Preprint.

[19] I. Z. Ruzsa. Generalized arithmetical progressions and sumsets, Acta Math. Hungar. 65 4 (1994), 379–388.

[20] T. Sanders. On the bogolyubov-ruzsa lemma. arXiv:1011.0107.

[21] Y. Shalom and T. Tao, A finitary version of Gromov's polynomial growth theorem. arXiv:0910.4148

[22] A. M. Turing, Finite approximations to Lie groups. Annals of Math. 39 1 (1937), 105–111.